

Transformative Design Catapult



Sandia National Laboratories

DARPA-BAA-16-39

Annual Report: Calendar Year 2017

Proposal Title	Transformative Design Catapult
Technical Area	2
Focus Area(s), if applicable	NA
Proposer Organization	Sandia National Laboratories
Type of Organization	FFRDC
Proposer Reference Number, if any	
Technical Point of Contact (POC)	Name: Miguel Aguilo Address: PO Box 5800 MS 0897 Albuquerque, NM, 87122 Telephone: 505-845-8964 Email: maguilo@sandia.gov
Administrative POC	Name: Susie Trujillo Address: PO Box 5800 MS 0845 Albuquerque, NM, 87122 Telephone: 505-284-4292 Email: struji@sandia.gov
Award Instrument Requested	WFO reimbursable contract
Budget	Total: \$1,750,000
Place(s) of Performance	Sandia National Laboratories
Period of Performance	4 Years
Other Team Members (subcontractors and consultants), if any	Technical POC Name: Brett Clark Organization: Sandia National Laboratories Organization Type: FFRDC Technical POC Name: Joshua Robbins Organization: Sandia National Laboratories Organization Type: FFRDC Technical POC Name: Paul Stallings Organization: Sandia National Laboratories Organization Type: FFRDC
Date Report was Prepared	Friday, January 5, 2018
Proposal Validity Period (minimum 120 days)	4 years

TABLE OF CONTENTS

ACRONYMS	4
1. EXECUTIVE SUMMARY.....	5
2. DEVELOPMENT PLAN SUMMARY	6
3. ACCOMPLISHMENTS	7
3.1 INTEROPERABLE DESIGN PLATFORM.....	7
3.2 MODELING.....	9
3.3 SYNTHESIS.....	11
3.4 ANALYSIS.....	12
3.5 TECHNICAL ASSISTANCE	15
4. FINANCIAL STATUS	16
5. BARRIERS	17
6. ACKNOWLEDGMENTS	17

ACRONYMS

API: Application Programming Interface

CPU: Central Processing Unit

CT: Computerized Axial Tomography

GCMMA: Globally Convergent Method of Moving Asymptotes

GPU: Graphics Processing Unit

HPC: High Performance Computing

LLNL: Lawrence Livermore National Laboratories

MIMD: Multiple Instruction Streams, Multiple Data Streams

MMA: Method of Moving Asymptotes

MPMD: Multiple Program Multiple Data

OC: Optimality Criteria

ROL: Rapid Optimization Library

SIMD: Single Instruction Stream, Multiple Data Streams

SIMT: Single Instruction Stream, Multiple Threads

SISD: Single Instruction Stream, Single Data Stream

SNL: Sandia National Laboratories

SPMD: Single Program Multiple Data

TRAL: Trust Region Augmented Lagrangian

TRBC: Trust Region Bound Constrained

UC – Boulder: University of Colorado Boulder

1. EXECUTIVE SUMMARY

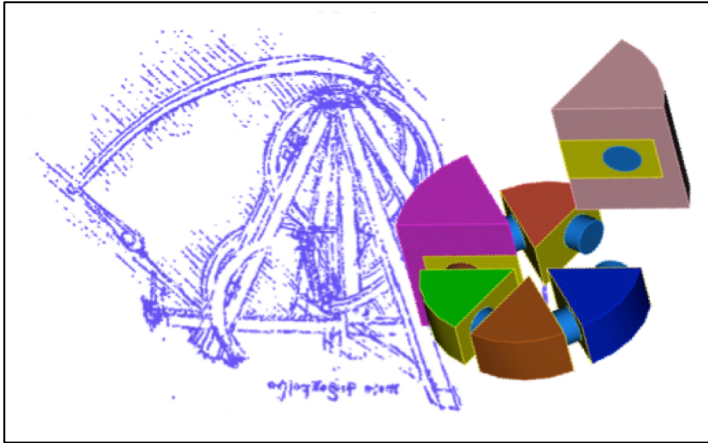


FIGURE 1: THE MAIN CONCEPT IS TO PROVIDE A COMMON, INTER-OPERABLE DESIGN PLATFORM WITH THE GOAL OF CATAPULTING RESEARCH AND DEVELOPMENT EFFORTS.

During calendar year 2017, Sandia National Laboratories (SNL) made strides towards developing an open portable design platform rich in high-performance computing (HPC) enabled modeling, analysis and synthesis tools. The main focus was to lay the foundations of the core interfaces that will enable plug-n-play insertion of synthesis optimization technologies in the areas of modeling, analysis and synthesis. SNL achieved the following milestones

during calendar year 2017:

1. Provided personal, systematic and sustained technical assistance;
2. Supplied existing software technology; and
3. Defined milestone tests.

In addition, SNL embarked on new challenges:

1. Development of an interoperable and portable HPC design platform;
2. Integration of technologies developed on the DARPA EQUiPS program; and
3. Development of an optimization aware analysis tool portable across multiple platforms.

Transforming current design practices is complex, multi-faceted and monumental in scope. During calendar year 2018, SNL will continue to provide personal, systematic and sustained technical assistance and available software tools for performers to leverage, augment and experiment with. The measures of success will be to:

1. Provide a common, interoperable design platform and catapult research and development activities (Figure 1);
2. Eliminate the mundane and unnecessary tasks of developing existing infrastructures and software tools; and
3. Provide personal, systematic and sustained technical support.

SNL is certain that by committing to these principles we can assure the success of the TRADES program.

TASK 1 (T1)	Type	Org.	Meshing Support	Q1, 2017	Q3, 2018
-T1.1	RG	SNL	Define PLATO-Engine API for in-situ validation analysis	■	Q1, 2017
-T1.2	CD	SNL	Develop PLATO-Engine API for in-situ validation analysis	■	Q4, 2017
-T1.3	TD	SNL	Demonstrate workflow for automatic in-situ validation analysis	▨	Q3, 2018
-T1.4	SR	SNL	Release PLATO-Engine meshing API and support growth		Q1, 2019
TASK 2 (T2)	Type	Org.	Synthesis Support	Q1, 2017	Q1, 2019
-T2.1	RG	SNL	Define generic PLATO-Engine synthesis API	■	Q1, 2017
-T2.2	CD	SNL	Develop generic PLATO-Engine synthesis API	■	Q4, 2017
-T2.3	TD	SNL	Demonstrate synthesis optimization workflow	■	Q4, 2018
-T2.4	SR	SNL	Release PLATO-Engine synthesis API and support growth		Q1, 2019
TASK 3 (T3)	Type	Org.	Analysis Support	Q1, 2017	Q4, 2018
-T3.1	RG	SNL	Define generic PLATO-Engine analysis API	■	Q1, 2017
-T3.2	CD	SNL	Develop generic PLATO-Engine analysis API	■	Q4, 2017
-T3.3	TD	SNL	Demonstrate analysis workflow	■	Q4, 2018
-T3.4	SR	SNL	Release PLATO-Engine analysis API and support growth		Q1, 2019
TASK 4 (T4)	Type	Org.	Reporting	Q1, 2017	Q4, 2020
-T4.1	RP	SNL	Annual report	■	Q4, 2017
-T4.2	RP	SNL	Annual report		Q4, 2018
-T4.3	RP	SNL	Annual report		Q4, 2019
-T4.4	RP	SNL	Final report		Q4, 2020

Completed
 In Progress

FIGURE 2: RESEARCH AND DEVELOPMENT PLAN.

2. DEVELOPMENT PLAN SUMMARY

Figure 2 shows the research and development plan proposed by SNL to DARPA. The ensuing list numerates the deliverables for calendar year 2017:

1. **Tasks 1.1 & 1.2:** Defined and developed interoperable application programming interfaces (API) for hands-free validation/verification analysis.
2. **Tasks 2.1 & 2.2:** Defined and developed interoperable API for HPC synthesis optimization.
3. **Tasks 3.1 & 3.2:** Defined and developed interoperable API for HPC analysis during synthesis optimization.

The aforementioned deliverables for calendar year 2017 were achieved. Figure 2 also shows that SNL met these deliverables related to the design platform, i.e. Plato Engine:

1. **Task 2.3:** Demonstrate interoperability and workflow with multiple gradient-based synthesis optimization algorithms.
2. **Task 3.3:** Demonstrated interoperability and workflow with multiple analysis tools.

In addition, SNL and its partners at University of Colorado Boulder (UC-Boulder), developed a new HPC model generation toolkit capable of automatically creating complex geometric models

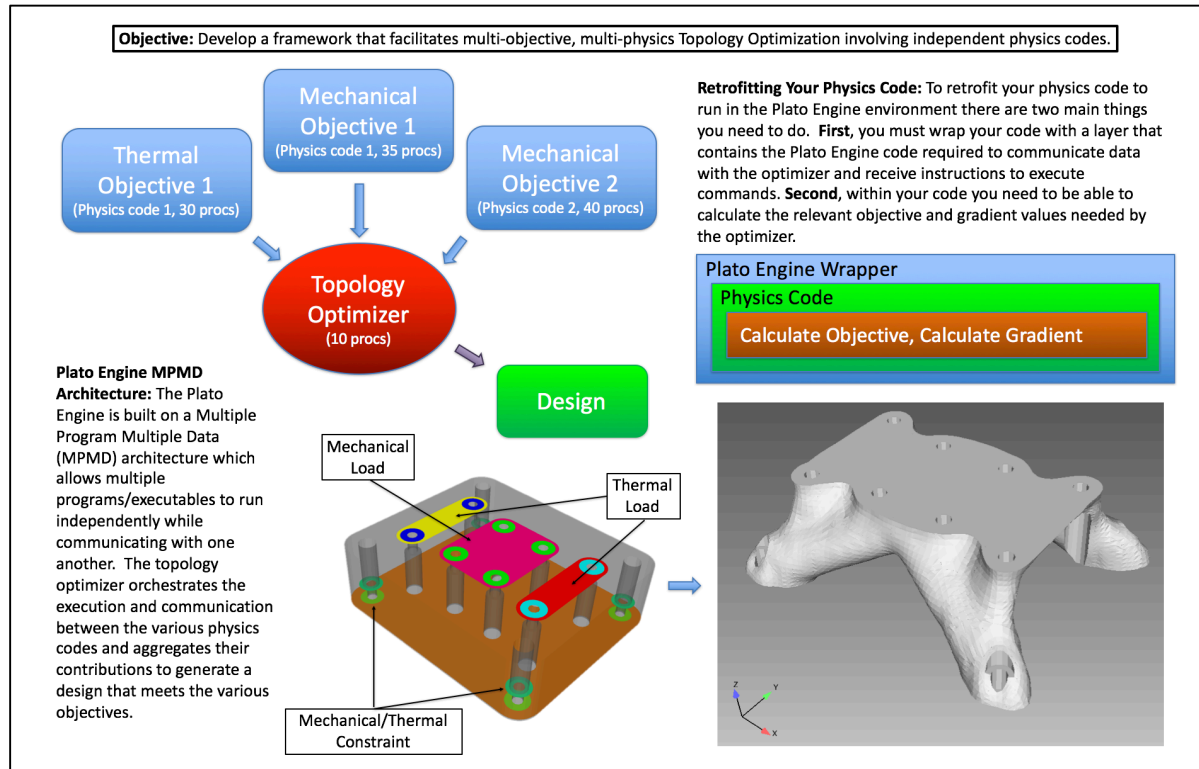


FIGURE 3: EXAMPLE SYNTHESIS OPTIMIZATION PROBLEM SOLVED WITH THE INTEROPERABLE DESIGN PLATFORM USING AN MPMD PROGRAMMING MODEL.

from level set, or voxel, representations. This toolkit was demonstrated on a free-form synthesis optimization problem, i.e. topology optimization problem, and on the femur challenge problem. The latter example used the computerized axial tomography (CT) scan data provided by DARPA to reconstruct a geometric model of the femur. SNL and UC-Boulder are currently testing the interfaces to ensure that the toolkit performs optimally when released.

3. ACCOMPLISHMENTS

3.1 INTEROPERABLE DESIGN PLATFORM

The main development goal for calendar year 2017 was to lay the foundations of the HPC design platform. This entailed careful crafting of lightweight interfaces that enable intercommunication of modeling, analysis and synthesis data using a multiple program, multiple data (MPMD) parallel programming model. This MPMD model allows multiple programs/executables to run independently while communicating with one another. The synthesis optimization algorithm orchestrates the execution and communication between the various analysis codes and aggregates their contributions to generate a design that meets multiple design criteria. This HPC architecture, among other features, facilitates the following:

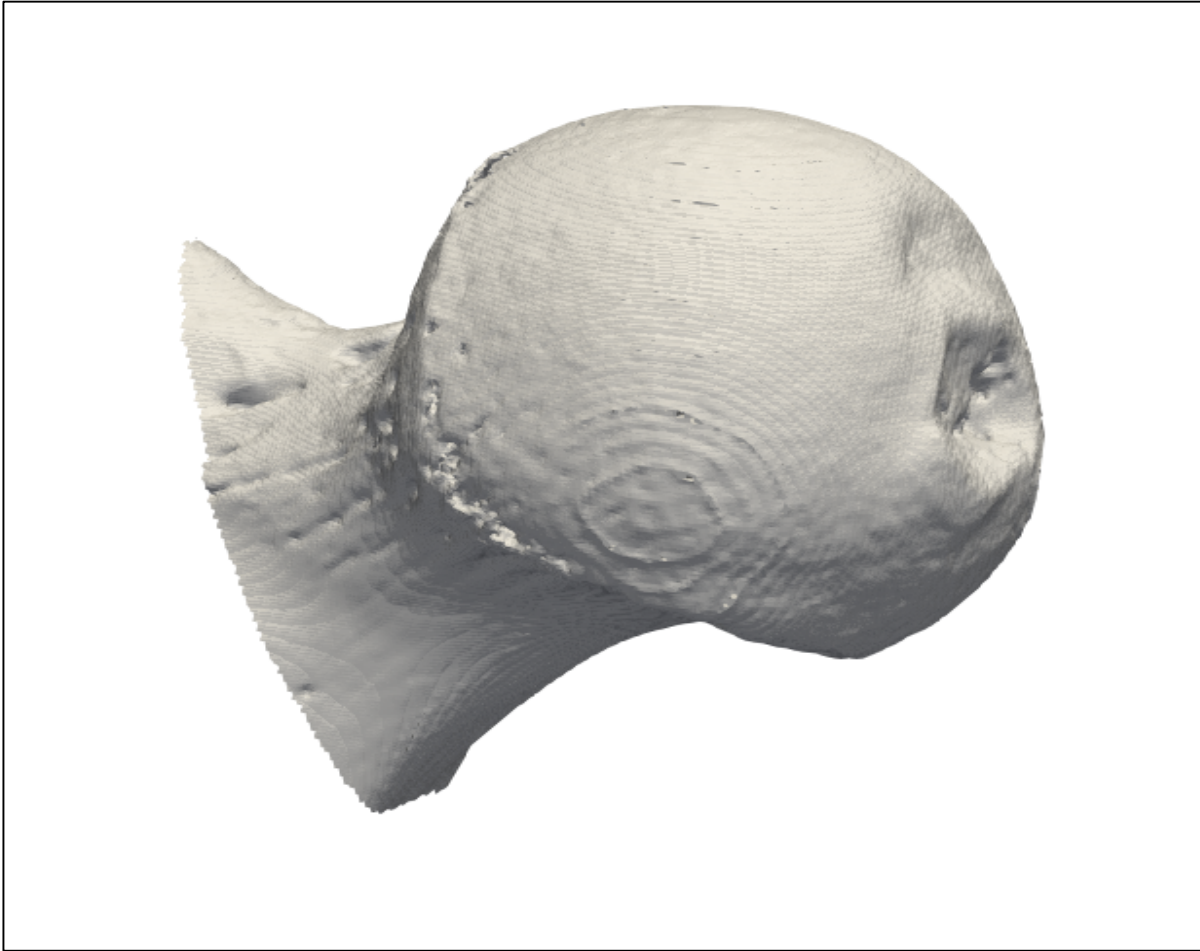


FIGURE 4: ONE QUARTER OF THE FEMUR MODEL CONSTRUCTED FROM THE CT SCAN DATA PROVIDED BY DARPA TRADES.

1. Combination of multiple analysis tools for multi-physics synthesis optimization;
2. Simultaneous analysis of multiple loads during synthesis optimization; and
3. Efficient solution of synthesis optimization problems with uncertain model parameters.

This design platform enables the intercommunication of modeling, analysis and synthesis data using a single program, multiple data (SPMD) parallel programming model. Therefore, if an application requires modeling, analysis and synthesis tools to share the same number of processors and avoid communication overhead, the Plato Engine platform will facilitate users to run their synthesis optimization problem with an SPMD model. Therefore, once users implement the five functions that define the Plato Engine interface, MPMD and SPMD parallel programming models will be supported. SNL plans to test the SPMD parallel programming model during calendar year 2018 to ensure that the Plato Engine interface properly supports both parallel programming models.

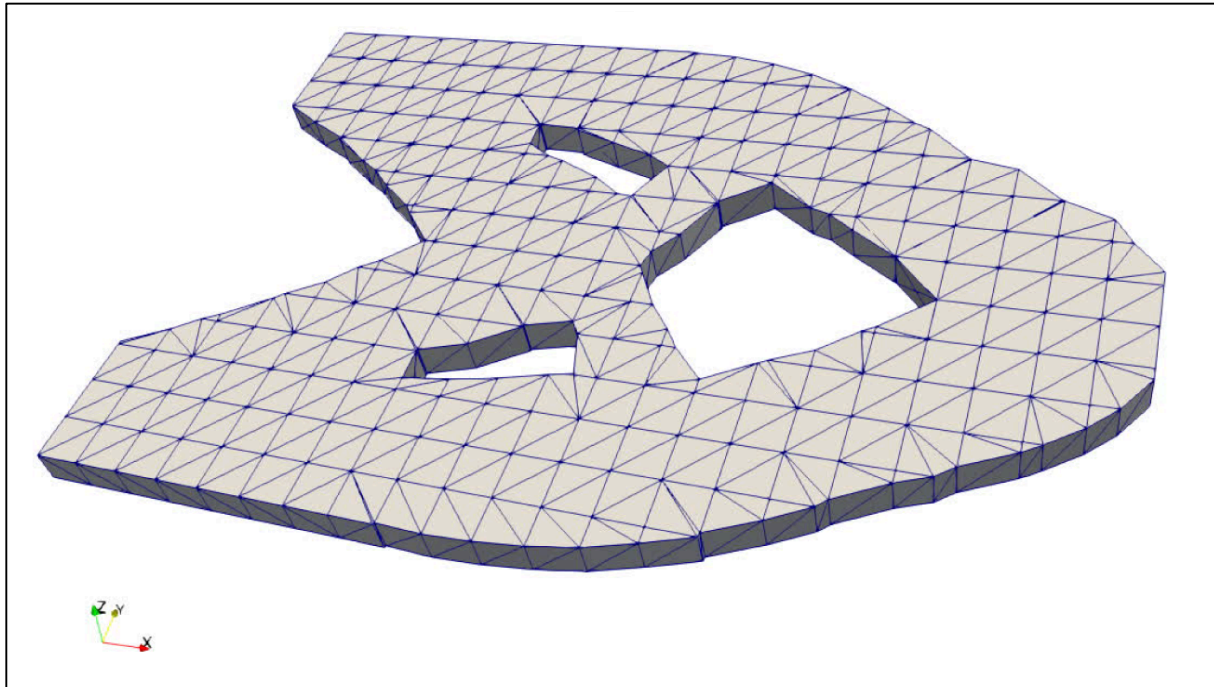


FIGURE 5: OPTIMAL GEOMETRY COMPUTED FROM A FREE-FORM (I.E. TOPOLOGY OPTIMIZATION) SYNTHESIS OPTIMIZATION PROCESS.

SNL also thinks that the new Plato Engine architecture could support cloud computing. The key is to develop the data communication layer that enables Plato to transfer and receive data from the modeling, analysis and synthesis tools running in the cloud. The existing data communication layer will be used to support SPMD and MPMD parallel programming models. During calendar year 2018 SNL could also explore how this same data communication layer could be used to support cloud computing. The interfaces of SNL modeling, analysis and synthesis tools do not need major modifications; thus, SNL would leverage these tools to test the cloud computing paradigm.

Lastly, unit test and regression test suites were developed to ensure that the design platform is nightly tested with multiple synthesis optimization algorithms and analysis tools (Sierra Structural Dynamics, Albany, etc.). The outlook for calendar year 2018 is the integration of more synthesis and analysis tools into the Plato ecosystem. Indeed, the Rapid Optimization Library (ROL), the electrostatic code Aleph, the graphic processing unit (GPU) accelerated multi-physics solver Alexa and Sierra Thermo-Mechanics & Fluids Module should be integrated into the Plato ecosystem. Furthermore, SNL and Lawrence Livermore National Laboratories (LLNL) are exploring a partnership that will enable LLNL to integrate their electromagnetics solver into the Plato ecosystem during calendar year 2018.

3.2 MODELING

The modeling accomplishments for calendar year 2017 are summarized below:

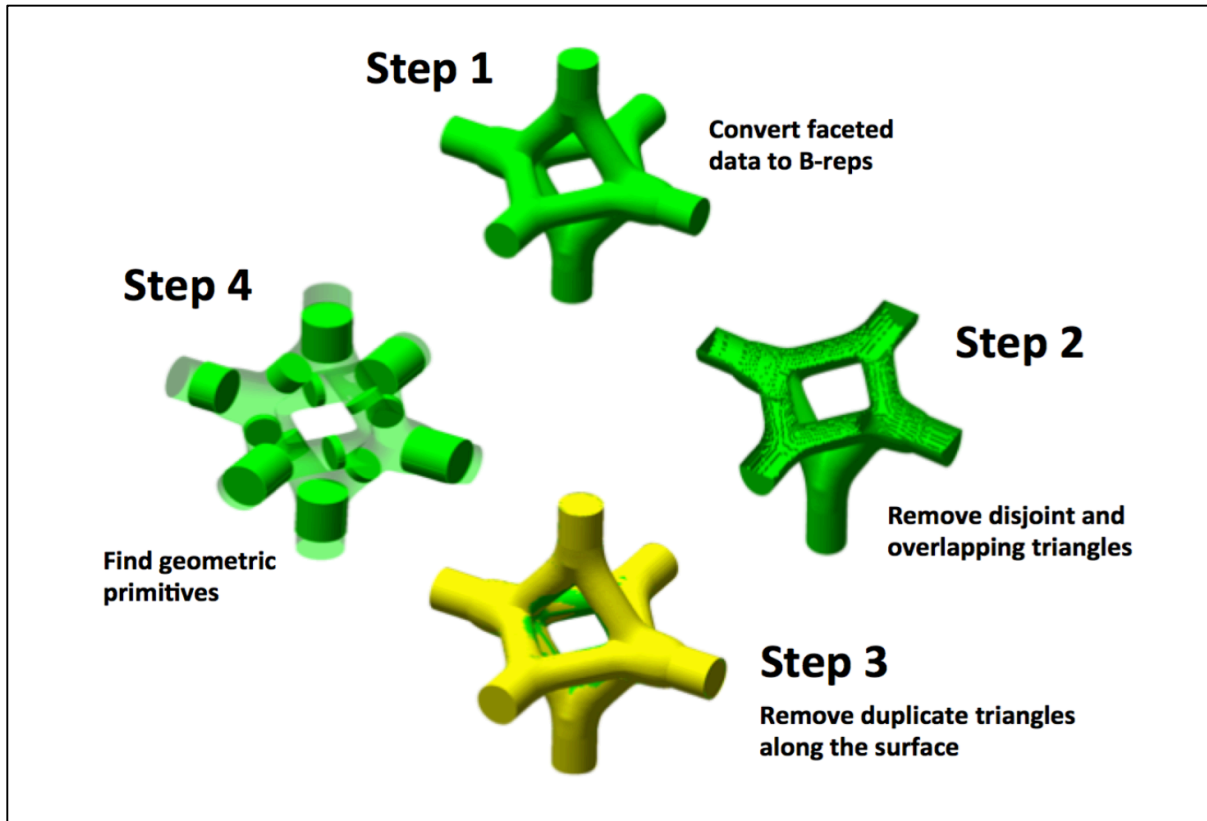


FIGURE 6: PROCESS TO AUTOMATICALLY IDENTIFY COMMON GEOMETRIC PRIMITIVES. FIRST, CONVERT FACETED DATA TO BOUNDARY REPRESENTATION (B-REPS). SECOND, REMOVE ANY DISJOINT AND OVERLAPPING TRIANGLES. THIRD, REMOVE DUPLICATE TRIANGLES ON THE BOUNDARY. FINALLY, IDENTIFY THE GEOMETRIC PRIMITIVES.

1. **Model Generation Toolkit¹:** Defined and developed APIs to enable hands-free verification simulations. Complex geometric models are created from a level set, or voxel, representation of the geometry. The model generation toolkit supports multiple material representations and is designed for HPC, which enables the hands-free creation of large-scale geometric models that are impossible to create without HPC. Figure 4 shows a femur model that was built from CT scan data. The toolkit took approximately four minutes on twelve processors to create the full femur model. The toolkit is also impacting work at SNL. Indeed, the toolkit will be used to complete a major design milestone during calendar year 2018. Figure 5 shows how SNL applied this toolkit to solve a free-form synthesis optimization problem and demonstrate the technology. By using a level set representation of the geometry, the toolkit constructed the optimal geometric model used for analysis.
2. **Primitive Identification Toolkit:** Developed a stand-alone, C++ library that automatically identifies common geometric features from faceted data, see Figure 6. This toolkit will be used by SNL and UC-Boulder to automatically identify common primitives that emerge

¹This work is a partnership between University of Colorado Boulder and Sandia National Laboratories.

during free-form synthesis optimization. This feature enables the optimizer to work with less number of design variables and local sensitivity calculations as the synthesis optimization problem progresses, making the optimizer faster and more intelligent. The primitive identification toolkit also permits the parameterization of level set representations, which enables parametric optimization (i.e. shape optimization). SNL plans to support more primitives during calendar year 2018.

3.3 SYNTHESIS

The synthesis optimization libraries developed during calendar year 2017 were designed to support multiple computer architectures as well as parallel programming models. Indeed, the ensuing list highlights the computer architectures supported:

1. **Single Instruction Stream, Single Data Stream (SISD):** Supports no parallelism.
2. **Single Instruction Stream, Multiple Data Streams (SIMD):** Supports multiple data streams in a single instruction stream, e.g. one GPU.
3. **Multiple Instruction Streams, Multiple Data Streams (MIMD):** Supports multiple autonomous processors simultaneously performing multiple instructions on distinct data streams, e.g. multicore processors and distributed architectures using either one shared or distributed memory space.
4. **Single Instruction Stream, Multiple Threads (SIMT):** Supports combining a single instruction stream, multiple data stream model with multithreading.
5. **Cloud Computing:** Enables access to shared pools of computing devices.

In addition to the aforementioned computer architectures, the optimization libraries support MPMD and SPMD parallel programming models. These features display the level of interoperability that these optimization libraries support. The ensuing list highlights the features offered by these libraries:

1. **Kernel Filter Library:** Stand-alone C++ kernel filter library used to smooth the geometry by penalizing rough surfaces, control the allowable feature size and stabilize the synthesis optimization formulation. The library also features multiple parallel search algorithms that can be used outside the kernel filter. Performers can start leveraging and using this tool by implementing the five functions that define the Plato-Engine interface, see Section 3.1.
2. **Gradient-Based Optimization Algorithms:** Stand-alone C++ optimization library suited for HPC. Several optimization algorithms were tightly integrated with Sierra Structural Dynamics, which prevented other analysis codes from leveraging the synthesis optimization algorithms. These algorithms were removed from Sierra Structural Dynamics Module and interoperable interfaces were defined and developed to enable other analysis codes to use the algorithms. The ensuing list highlights some of the algorithms available in the library:

1. Optimality Criteria (OC): First-order (i.e. supports gradient but no Hessian information) optimization algorithm tailored for compliance minimization problems with bound constraints and single or multiple linear inequality constraints. The algorithm supports multiple materials, i.e. design variable fields, suitable for multi-material synthesis optimization problems with bound constraints and single/multiple linear inequality constraints. The algorithm does not support a globalization strategy, making it prone to local minima.
2. Method of Moving Asymptotes (MMA): First-order nonlinear programming optimization algorithm that supports bound, linear and nonlinear inequality constraints and multiple design variable fields. Compared to the OC algorithm, it can be applied to more general nonlinear programming optimization problems. Therefore, the MMA algorithm is not tailored for compliance minimization problems. However, similar to the OC algorithm, the algorithm is prone to local minima since it does not support a globalization strategy.
3. Globally Convergent Method of Moving Asymptotes (GCMMA): First-order nonlinear programming optimization algorithm that supports bound, linear and nonlinear inequality constraints and multiple design variable fields. Similar to the MMA algorithm, it is suited for more general nonlinear programming optimization problems. However, contrary to the MMA algorithm, the GCMMA algorithm is less prone to local minima since it supports a globalization strategy.
4. Trust Region Bound Constrained (TRBC): Matrix-free second-order nonlinear programming optimization algorithm that supports gradient and Hessian information. If the user provides Hessian information, the algorithm is capable of achieving quadratic convergence rates. Contrary, first-order optimization algorithms can only achieve linear convergence rates. If the user provides no Hessian information, the algorithm supports multiple features to approximate the Hessian information with minimal computational cost. The trust region scheme is a theoretically sound globalization strategy that enables the algorithm to be less prone to local minima. However, the only downfall of the TRBC algorithm is that it only supports bound constraints. Hence, it is not suited for synthesis optimization problems with inequality constraints.
5. Trust Region Augmented Lagrangian (TRAL): Matrix-free second-order nonlinear programming optimization algorithm that has all the features supported by the TRBC algorithm. However, contrary to the TRBC algorithm, the TRAL algorithm supports linear and nonlinear inequality constraints.

3.4 ANALYSIS

During calendar year 2017 SNL embarked on a new research challenge: the development of an optimization aware, multi-physics solver that is portable across multiple computer architectures. The flexibility achieved with this templetized multi-physics solver is such that it

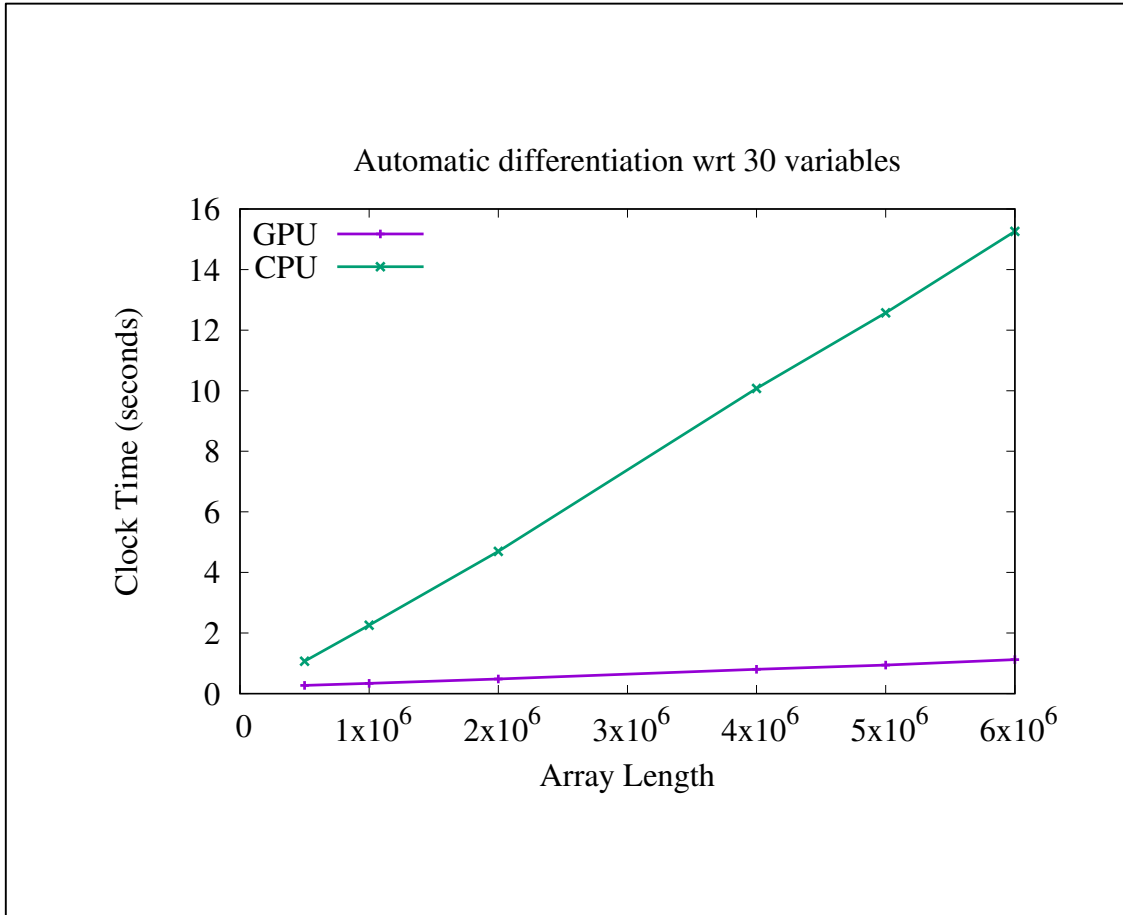


FIGURE 7: THIS PLOT SHOWS THE SPEEDUPS ACHIEVED WITH THE GPU, VERSUS THE CPU, USING AUTOMATIC DIFFERENTIATION TO COMPUTE THE SENSITIVITY OF A CRITERIA WITH RESPECT TO A SET OF DESIGN VARIABLES. IN THIS EXAMPLE, EACH NODE IN THE COMPUTATIONAL MESH DEPENDED ON 30 DESIGN VARIABLES. HENCE, THE TOTAL NUMBER OF DESIGN VARIABLES EQUALS THE ARRAY LENGTH TIMES 30.

facilitates the use of different floating-point types during analysis. For instance, SNL is assisting Etaphase compare Unum and IEEE floating-point types. This exercise would have been difficult, if not impossible to be performed, without the templetized analysis tool. Furthermore, since the solver was designed to be portable across multiple computer architectures, the solver supports SISD, SIMD, MIMD and SIMT data parallelism and MPMD and SPMD parallel programming models.

SNL also added other features that enhance the functionalities available in the multi-physics solver. The ensuing list highlights some of these features:

1. **Automatic Differentiation:** Enables automatic computation of the sensitivities needed to solve synthesis optimization problems. For instance, Figure 7 shows the central processing unit (CPU) and GPU time that the automatic differentiation tool took to compute the sensitivities of a criterion defined as the measurement of the misfit between a set of

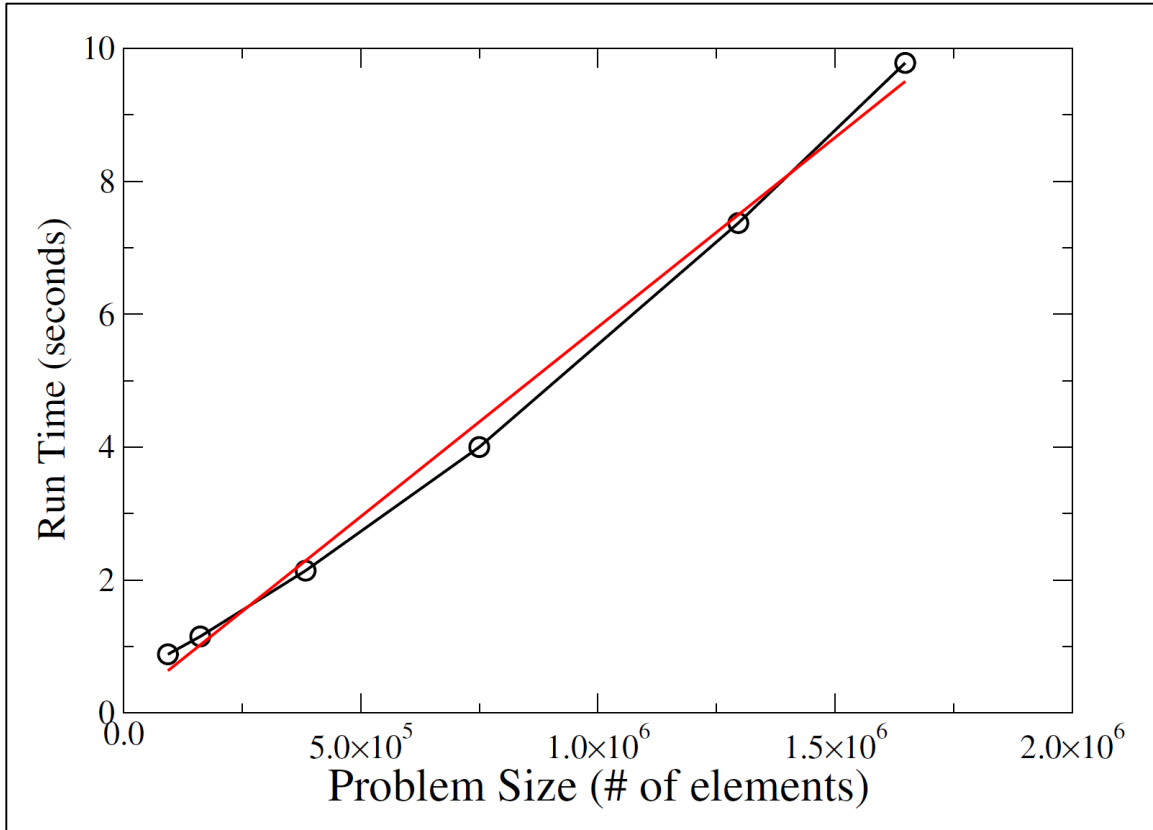


FIGURE 8: RUN TIMES FOR LINEAR STATICS SOLUTION ON A LAPTOP WITH A SINGLE NVIDIA QUADRO M4000M GPU. THE SOLVER CONVERGENCE WAS SET TO $1E-12$ OF INITIAL RESIDUAL. THE RED LINE DENOTES THE THORETICAL SCALING AND THE BLACK LINE DENOTES THE ACTUAL SCALING.

experimental measurements and simulated state solution with respect to a set of design variables. Figure 7 not only highlights the portability and interoperability of the automatic differentiation tool, but it also highlights the efficiency of the GPU. This tool will allow developers to avoid the tedious task of deriving and coding sensitivities and focus on improving the algorithms needed to solve synthesis optimization problems.

2. **Adaptive Mesh Refinement:** Enables adaptive mesh refinement of tetrahedron and triangle meshes, with a focus on scalable HPC performance. It is intended to provide adaptive functionality to existing simulation codes. This tool allows users to reduce discretization error and number of degrees of freedom during analysis and enable modeling of moving objects and evolving geometries.
3. **Speedups:** Figure 7 shows the speedups that the GPU version of the automatic differentiation tool enables over its CPU counterpart. Indeed, the GPU-accelerated solver was able to solve a $2e6$ degree of freedom heat conduction problem in approximately 0.8 seconds using a single NVIDIA Tesla P100 GPU. Figure 8 shows an additional example highlighting the performance of the GPU-accelerated linear static module in the multi-physics analysis tool. It is clear that the GPU is performing as desired. SNL is excited about

the multiple possibilities that this solver offers and will continue maturing the analysis tool during calendar year 2018.

3.5 TECHNICAL ASSISTANCE

SNL fulfilled its commitment to provide performers with technical assistance during calendar year 2017. SNL contacted all TA1 performers, gathered technical and software needs and explored multiple partnerships. The ensuing list highlights some of these interactions:

1. **Technical Outreach:** SNL was able to assist multiple partners in academia:
 1. Prof. Kurt Maute: SNL and Prof. Maute are developing the model generation toolkit that can automatically generate a geometry model from a level set representation. SNL plans to integrate this tool into the Plato ecosystem and make it available to performers during calendar year 2018.
 2. Prof. Carolyn Seepersad: A Plato license was approved and a copy of the Plato software was sent to Prof. Seepersad during calendar year 2017. Prof. Seepersad is expected to use Plato for her machine learning research work.
 3. Prof. Jorg Peters: Prof. Peters expects to visit SNL during calendar year 2018 and assist SNL with the integration of his modeling software into the Plato ecosystem. This interaction will allow Prof. Peters to test his modeling algorithms with SNL's analysis and synthesis tools.
 4. Prof. Matt Campbell: Multiple teleconferences were held to discuss the synthesis optimization algorithms available in Plato. SNL encourage Prof. Campbell to leverage ROL, an open source optimization library developed by SNL.
 5. Prof. Vadim Shapiro: Multiple teleconferences were held to discuss interoperability. SNL and Prof. Shapiro shared and discussed multiple approaches that enhanced the interoperability of the Plato Engine platform.
2. **Etaphase**: SNL provided sustained technical assistance to Etaphase during calendar year 2017. Indeed, SNL and Etaphase are partnering to test the Unum floating-point type with the templetized GPU-accelerated solver.
3. **Siemens**: Siemens plans to use Albany as their main analysis platform during the TRADES program. SNL, after multiple teleconferences, aided Siemens understanding of Albany's infrastructure and workflow. These discussions allowed Siemens to independently begin developing a multiscale material model in Albany.
4. **Lawrence Livermore National Laboratories**: SNL and LLNL are pursuing a partnership that will enable LLNL to leverage the modeling, analysis and synthesis tools in Plato. The goal is to integrate LLNL electromagnetics solver with Plato during calendar year 2018. SNL will assist LLNL implement the five functions that define the Plato Engine interface.

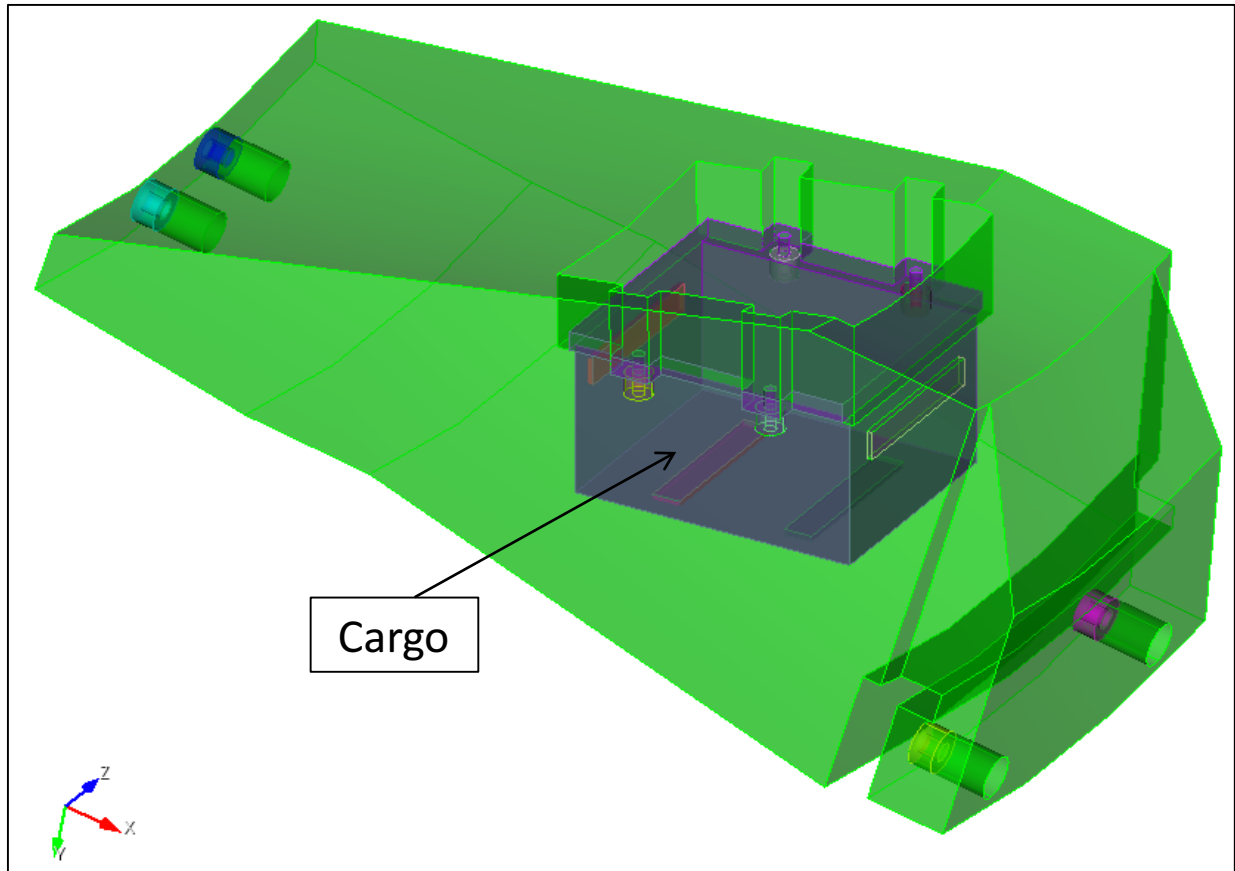


FIGURE 9: BRACKET GEOMETRY

5. **United Technologies:** SNL developed a stand-alone C++ library that takes a level set representation of the geometry and automatically detects tubes and outputs the parameters that define these tubes. This tool enable United Technologies to accelerate their machine learning research work.
6. **Challenge Problems:** SNL provided one of the two synthesis challenge problems presented by DARPA to the performers. The challenge problem aims to design a lightweight bracket that preserves the structural integrity of the cargo shown on Figure 9. The bracket will be subjected to multiple loading configurations and must preserve its structural integrity for all loading combinations. In addition, the first natural frequency must be maximized to avoid having the cargo vibrate once connected to a next level assembly. SNL also assisted DARPA on the second synthesis challenge problem by providing a concise report that depicted all the loading configurations applied to the racecar bracket.

4. FINANCIAL STATUS

See financial report submitted with the annual report.

5. BARRIERS

No major technical barriers to report.

6. ACKNOWLEDGMENTS

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.